Experimentally Investigating Geometric Locality in Deep Neural Networks

Abstract

Deep neural networks are commonly analysed in terms of expressivity, approximation theory and optimization. This note suggests a complementary viewpoint based upon **geometric locality**. The central idea is that successful learning may depend not merely upon preserving topology, but upon preserving local geometric relationships throughout the computational pipeline. If neighbouring data points remain geometrically related after many layers of computation, statistical regularities may remain accessible to gradient-based learning. Fortunately, this hypothesis leads naturally to a number of experimentally accessible questions that can be investigated using existing neural network architectures.

From Topology to Geometric Locality

Topology asks qualitative questions.

- Is the space still connected?
- Has it been torn?
- Have separate regions been glued together?

These are important mathematical properties.

However, they may not be sufficient for understanding why deep neural networks learn successfully.

A smooth mapping can preserve topology while still stretching nearby points over enormous distances.

From the perspective of statistical learning, such a transformation may be highly undesirable.

A stronger requirement may be **geometric locality**.

Informally, geometric locality means that

- nearby inputs generally remain nearby,
- neighbourhoods remain coherent,
- local distortions remain bounded,
- geometry changes smoothly as one moves through the representation.

This is a geometric rather than purely topological property.

The hypothesis is therefore:

> **Successful neural networks preserve sufficient geometric locality throughout computation that statistical relationships in natural data remain exploitable after many layers.**

This hypothesis immediately suggests a number of measurable experiments.

---

1. Local Expansion

Measure how rapidly nearby points separate.

For neighbouring samples

$$x, \qquad x + \epsilon$$

measure

$$\frac{\|f(x+\epsilon) - f(x)\|}{\|\epsilon\|}.$$

Questions:

- Does this remain bounded throughout training?
- Does it differ between successful and unsuccessful architectures?
- Does it correlate with generalization?

---

2. Neighbourhood Preservation

For each sample

- compute its nearest neighbours before a layer,
- compute its nearest neighbours after the layer.

Measure how many neighbours remain neighbours.

Questions:

- Is neighbourhood preservation maintained across depth?
- Does preservation improve during training?
- Are deeper representations surprisingly stable?

---

3. Context Continuity

For architectures based upon context-selected linear mappings, or for ReLU activation regions, compare neighbouring local linear models.

Measure quantities such as

$$\|A_i - A_j\|.$$

Questions:

- Do neighbouring contexts converge toward similar mappings?
- Does backpropagation implicitly smooth neighbouring models?
- Does context continuity increase during training?

---

4. Jacobian Smoothness

Measure how rapidly local derivatives change.

Compare

$$J(x)$$

with

$$J(x+\epsilon).$$

Questions:

- Does Jacobian variation decrease during optimization?
- Are successful networks locally smoother than expected?
- Do different architectures exhibit characteristic Jacobian behaviour?

---

5. Activation Stability

For ReLU networks, apply tiny perturbations to inputs.

Measure

- the number of ReLU state changes,
- the fraction of neurons that change activation.

Questions:

- Do trained networks become activation-stable?
- Are switching events relatively sparse?
- Does stability improve during learning?

---

6. Geometric Dispersion

Track small clouds of nearby points.

Measure

- cloud diameter,
- covariance,
- principal directions,
- local volume.

Questions:

- Do neighbourhoods remain compact?
- Which layers contribute most to geometric distortion?
- Is distortion anisotropic?

---

7. Layerwise Trajectory Analysis

Move continuously through input space.

Observe how internal representations evolve.

Questions:

- Are trajectories smooth?
- Where do abrupt changes occur?
- Do trajectories become smoother as training progresses?

---

8. Atlas Continuity

Construct the graph of local linear regions.

Connect neighbouring regions differing by small context changes.

Measure

- matrix differences,
- Jacobian differences,
- continuity across region boundaries.

Questions:

- Does the network form a coherent atlas?
- Are neighbouring regions statistically related?
- Does optimization encourage compatibility?

Measuring During Training

Perhaps the most interesting opportunity is to perform these measurements **throughout optimization rather than only after training has completed.**

This raises several important questions.

- Does geometric locality emerge gradually?
- Does it appear suddenly during training?
- Does it develop before accuracy improves?
- Does improved geometric locality predict improved generalization?
- Do different optimizers produce different geometric behaviour?
- What effect do weight decay, normalization and residual connections have on locality?

Viewing optimization as the evolution of computational geometry may provide insights unavailable from loss curves alone.

A Different Perspective on Backpropagation

Backpropagation is usually interpreted as an optimization algorithm.

An alternative viewpoint is that it also constructs a geometrically coherent family of local computational models.

Each local linear mapping becomes

- statistically averaged,
- specialized to its local context,
- compatible with neighbouring contexts.

If correct, optimization is simultaneously learning a function and organizing the geometry through which that function is computed.

Why These Experiments Matter

Many theoretical analyses of deep networks concentrate on

- expressivity,
- approximation capability,
- optimization,
- parameter efficiency.

The experiments proposed here instead investigate the geometry of the learned computation itself.

Rather than asking

> "How complicated is the represented function?"

they ask

> "How well does the computation preserve local geometric structure while representing that function?"

This shift in viewpoint may provide a useful bridge between differential geometry, topology, switched systems, control theory and deep learning.

Whether geometric locality proves to be a fundamental principle or merely a useful descriptive tool is ultimately an experimental question. Fortunately, many of the measurements required to investigate this hypothesis are already accessible using existing neural network architectures and standard machine learning frameworks.